

Hidden Markov model tracking of continuous gravitational waves from a neutron star with wandering spin

S. Suvorova,^{1,2} L. Sun,^{2,*} A. Melatos,^{2,†} W. Moran,¹ and R. J. Evans³

¹*School of Electrical and Computer Engineering,
RMIT University, Melbourne, Victoria 3000, Australia*

²*School of Physics, University of Melbourne, Parkville, Victoria 3010, Australia*

³*Department of Electrical and Electronic Engineering,
University of Melbourne, Parkville, Victoria 3010, Australia*

(Dated: June 9, 2016)

Gravitational wave searches for continuous-wave signals from neutron stars are especially challenging when the star's spin frequency is unknown a priori from electromagnetic observations and wanders stochastically under the action of internal (e.g. superfluid or magnetospheric) or external (e.g. accretion) torques. It is shown that frequency tracking by hidden Markov model (HMM) methods can be combined with existing maximum likelihood coherent matched filters like the \mathcal{F} -statistic to surmount some of the challenges raised by spin wandering. Specifically it is found that, for an isolated, biaxial rotor whose spin frequency walks randomly, HMM tracking of the \mathcal{F} -statistic output from coherent segments with duration $T_{\text{drift}} = 10$ d over a total observation time of $T_{\text{obs}} = 1$ yr can detect signals with wave strains $h_0 > 2 \times 10^{-26}$ at a noise level characteristic of the Advanced Laser Interferometer Gravitational Wave Observatory (Advanced LIGO). For a biaxial rotor with randomly walking spin in a binary orbit, whose orbital period and semi-major axis are known approximately from electromagnetic observations, HMM tracking of the Bessel-weighted \mathcal{F} -statistic output can detect signals with $h_0 > 8 \times 10^{-26}$. An efficient, recursive, HMM solver based on the Viterbi algorithm is demonstrated, which requires $\sim 10^3$ CPU-hours for a typical, broadband (0.5-kHz) search for the low-mass X-ray binary Scorpius X-1, including generation of the relevant \mathcal{F} -statistic input. In a “realistic” observational scenario, Viterbi tracking successfully detects 41 out of 50 synthetic signals without spin wandering in Stage I of the Scorpius X-1 Mock Data Challenge convened by the LIGO Scientific Collaboration down to a wave strain of $h_0 = 1.1 \times 10^{-25}$, recovering the frequency with a root-mean-square accuracy of $\leq 4.3 \times 10^{-3}$ Hz.

PACS numbers: 95.85.Sz, 97.60.Jd

I. INTRODUCTION

Continuous-wave gravitational radiation from isolated and accreting neutron stars is a key target of long-baseline interferometers like the Laser Interferometer Gravitational Wave Observatory (LIGO) and Virgo detector [1]. Theory predicts that the signal is quasi-monochromatic. Emission occurs at simple rational multiples of the star's spin frequency f_* , for example f_* and $2f_*$ for mass quadrupole radiation from thermoelectric and magnetic mountains [2, 3], $4f_*/3$ for r-modes [4, 5], and f_* for current quadrupole radiation from non-axisymmetric flows in the neutron superfluid pinned to the stellar crust [6]. If the source exhibits electromagnetic pulsations, so that an ephemeris can be derived from absolute pulse numbering, i.e. $f_*(t)$ is known as a function of time t , it is customary to search for a signal using coherent matched filters like the maximum likelihood \mathcal{F} -statistic [7]. If an ephemeris is unavailable, coherent searches over multiple $f_*(t)$ templates indexed by the Taylor coefficients $f_*^{(k)}(0) = (d^k f_*/dt^k)_{t=0}$ become expensive computationally [8], and semi-coherent

methods are often preferred. Cross-correlation [9, 10], StackSlide [11], the Hough transform [12–14], PowerFlux [15–20] and TwoSpect [21, 22] are all examples of semi-coherent algorithms implemented by the LIGO Scientific Collaboration and applied to data from Science Runs 5 or 6.

Search methods that scan templates without guidance from a measured ephemeris are compromised if $f_*(t)$ wanders randomly. Radio and X-ray timing of pulsating neutron stars reveal that spin wandering is a widespread phenomenon. In isolated objects, it manifests itself as timing noise [23–25], exhibits a red Fourier spectrum with an auto-correlation time-scale of days to years [26, 27], and has been attributed variously to magnetospheric changes [28], superfluid dynamics in the stellar interior [27, 29–31], spin microjumps [32, 33], and fluctuations in the spin-down torque [34–36]. In accreting objects, spin wandering results from fluctuations in the magnetized accretion torque [37–39], due to transient accretion disk formation [40, 41] or disk-magnetospheric instabilities and reconnection events [42]. Again the auto-correlation time-scale is of the order of days [43], and fluctuations in $f_*(t)$ are accompanied by fluctuations in the X-ray flux [39]. Even when the ephemeris is measured electromagnetically, the gravitational-wave-emitting quadrupole may not be phase locked to the stellar crust, and wandering must still be accommodated [44].

* lings2@student.unimelb.edu.au

† amelatos@unimelb.edu.au

Hidden Markov model (HMM) methods offer one powerful strategy for detecting and tracking a wandering frequency [45]. The essential idea is to model $f_\star(t)$ probabilistically as a Markov chain of transitions between unobservable (“hidden”) frequency states and relate the hidden states to the observed data via a detection statistic. HMM frequency tracking enjoys a long track record of success in engineering applications ranging from radar and sonar analysis [46] to mobile telephony [47, 48]. It has been refined substantially since its introduction by Streit and Barrett [49] to include information about amplitude and phase [50] and embrace simultaneous tracking of multiple targets and frequencies [51, 52]. It delivers accurate estimation, when the signal-to-noise ratio (SNR) is low but the sample size is large [45], the situation normally confronting gravitational wave data analysis targeting continuous-wave sources.

In this paper, we implement and test a specific HMM scheme based on the classic Viterbi algorithm [45, 53]. The scheme is efficient and practical: its computational demands are modest, and it co-opts existing technology for LIGO continuous-wave searches based on the \mathcal{F} -statistic. The paper is organized as follows. In Section II, we formulate the search problem in HMM language and describe how to solve it with the Viterbi algorithm. In Section III, we apply Viterbi tracking to an isolated neutron star with $f_\star(t)$ walking randomly, using the \mathcal{F} -statistic to relate the hidden and observable states. We quantify the performance of the tracker as a function of signal strength using synthetic data. In Section IV we apply Viterbi tracking to a neutron star in a binary, again with $f_\star(t)$ wandering randomly, replacing the \mathcal{F} -statistic with a Bessel-weighted variant, and quantify its performance using synthetic data. Finally, to illustrate how the tracker performs in a “realistic” scenario, we apply it to the data set prepared for Stage I of the Scorpius X-1 (Sco X-1) Mock Data Challenge in Section V [54].

II. FREQUENCY TRACKING

HMM principles can be applied to gravitational wave frequency tracking in various ways, depending on the data format, the detection statistic, and any known constraints on the frequency evolution.

In this paper, we specialize to continuous-wave searches, where the raw data are packaged in 30-min short Fourier transforms (SFTs), during which $f_\star(t)$ remains localized within one Fourier bin. The SFTs are fed into a frequency-domain estimator $G(f)$, like the \mathcal{F} -statistic (isolated target) or Bessel-weighted \mathcal{F} -statistic (binary target). It is safe to assume that $f_\star(t)$ remains localized within one estimator frequency bin over a short enough time interval, even when $f_\star(t)$ walks randomly. Let T_{obs} be the total observation time, and let $T_{\text{SFT}} = 30 \text{ min}$ be the length of each SFT. For any particular astrophysical source, there exists an intermediate time-scale T_{drift} ($T_{\text{SFT}} < T_{\text{drift}} < T_{\text{obs}}$), over which

$f_\star(t)$ wanders by at most one estimator bin. For timing noise in isolated pulsars, radio timing experiments imply $T_{\text{drift}} \gtrsim \text{weeks}$ [23]. For accretion noise in low-mass X-ray binaries, such as Sco X-1, one can estimate $T_{\text{drift}} \approx 10 \text{ days}$ theoretically [54–56]. Our general strategy is to compute $G(f)$ for blocks of data of length T_{drift} then use the Viterbi algorithm to track peaks in $G(f)$ over the full observation interval $0 \leq t \leq T_{\text{obs}}$, effectively summing the estimator output semi-coherently by tracking $f_\star(t)$.¹ An important practical advantage of this approach is that it leverages the existing, efficient, thoroughly tested \mathcal{F} -statistic software infrastructure within the LIGO Algorithms Library (LAL), which is used extensively by the continuous-wave data analysis community [57].

In this section we formulate the problem as an HMM (Section II A–II C) and describe the recursive Viterbi algorithm for solving the problem (Section II D). We defer to future work the treatment of glitches, i.e. random, impulsive, spin-up events, where $f_\star(t)$ jumps over many estimator bins instantaneously [58, 59].

A. HMM formulation

An HMM is a probabilistic finite state automaton defined by a hidden (unobservable) state variable $q(t)$, which takes one of a finite set of values $\{q_1, \dots, q_{N_Q}\}$ at time t , and an observable state variable $o(t)$, which takes one of the values $\{o_1, \dots, o_{N_o}\}$. The automaton jumps between states at discrete times $\{t_0, \dots, t_{N_T}\}$. The jump probability from time t_n to time t_{n+1} depends only on the hidden state $q(t_n)$ at time t_n — the Markovian assumption — and is described by the transition probability matrix

$$A_{q_j q_i} = \Pr[q(t_{n+1}) = q_j | q(t_n) = q_i]. \quad (1)$$

The likelihood that the system is observed in state $o(t_n)$ at time t_n is described by the emission probability matrix

$$L_{o_j q_i} = \Pr[o(t_n) = o_j | q(t_n) = q_i]. \quad (2)$$

The model is completed by specifying the probability that the system occupies each hidden state initially, described by the prior vector

$$\Pi_{q_i} = \Pr[q(t_0) = q_i]. \quad (3)$$

Suppose that we observe the system transitioning through the sequence of observable states $O = \{o(t_0), \dots, o(t_{N_T})\}$. In general there exist $N_Q^{N_T+1}$ possible paths $Q = \{q(t_0), \dots, q(t_{N_T})\}$ through the hidden

¹ T_{drift} can be as short as T_{SFT} , the minimum time over which it makes sense to calculate $G(f)$, without harming the performance of the tracking algorithm. As the algorithm is semi-coherent, the strain sensitivity scales approximately as $(T_{\text{obs}}/T_{\text{drift}})^{1/4}$ [1].

states which are consistent with the observed sequence O . For a Markov process, the probability that the hidden path Q gives rise to the observed sequence O equals

$$P(Q|O) = L_{o(t_{N_T})q(t_{N_T})} A_{q(t_{N_T})q(t_{N_T-1})} \cdots L_{o(t_1)q(t_1)} \times A_{q(t_1)q(t_0)} \Pi_{q(t_0)}. \quad (4)$$

The most probable path Q^* is the one that maximizes $P(Q|O)$, i.e.,

$$Q^*(O) = \arg \max P(Q|O), \quad (5)$$

where $\arg \max(\dots)$ returns the argument that maximizes the function (\dots) . The Viterbi algorithm, presented in Section IID, provides a recursive, computationally efficient route to computing $Q^*(O)$ from (1)–(5). For computational reasons, we actually evaluate $\log P(Q|O)$, whereupon (4) becomes a sum of log likelihoods.

B. Frequency drift time-scale and jump probabilities

In our first application to isolated neutron stars (Section III), the hidden state variable is one-dimensional and equals the star's spin frequency at time t , i.e. $q(t) = f_*(t)$. Its allowed, discretized values correspond one-to-one to the frequency bins in the output of the frequency-domain estimator $G(f_*)$ computed over an interval of length T_{drift} , indexed by their central (say) frequencies f_{*i} , i.e. $q_i = f_{*i}$.² If the total search band covers the frequency range $f_{*,\min} \leq f_*(t) \leq f_{*,\max} = f_{*,\min} + B$, and if the bin width of $G(f_*)$ is $\Delta f_{\text{drift}} = 1/(2T_{\text{drift}})$, then the number of hidden states is given by $N_Q = B/\Delta f_{\text{drift}}$.

Radio and X-ray timing observations demonstrate that neutron star spin wandering is a continuous stochastic process in the absence of glitches [23, 24, 39]. In what follows we approximate it by an unbiased random walk or Wiener process. Choosing T_{drift} to satisfy

$$\left| \int_t^{t+T_{\text{drift}}} dt' \dot{f}_*(t') \right| < \Delta f_{\text{drift}} \quad (6)$$

guarantees that the transition probability matrix takes a simple tridiagonal form, with

$$A_{q_{i+1}q_i} = A_{q_iq_i} = A_{q_{i-1}q_i} = \frac{1}{3} \quad (7)$$

and all other entries zero. In other words, at each time step, $f_*(t)$ jumps at most one frequency bin up or down or stays in the same bin with equal probability $1/3$. Studies

show that the performance of a HMM tracking scheme is insensitive to the exact form of the probabilities $A_{q_jq_i}$, as long as they capture broadly the behaviour of the underlying jump process [45]. Equation (7) can be generalized to accommodate discontinuous glitches, but doing so lies outside the scope of this paper.³

In our second application to binary neutron stars (Sections IV and V), the hidden state variable is two-dimensional: we search over not only $f_*(t)$ but also the projected semimajor axis $a_0 = a \sin \iota$, where a is the semimajor axis of the orbit, and ι is the angle of inclination. In practice we divide the one-standard-deviation error bar on a_0 from electromagnetic measurements into N_{a_0} bins, whereupon the total number of hidden states equals $N_Q = N_{a_0} B / \Delta f_{\text{drift}}$. It is known astrophysically that a_0 does not change significantly over a typical observation ($T_{\text{obs}} \sim 1 \text{ yr}$), so the transition matrix is still given by equation (7), with the subscript i indexing the frequency bin f_{*i} but not a_0 . In other words, there is no dependence on a_0 in the two-dimensional version of equation (7). Therefore no separate subscript is needed to index a_0 .

C. Emission and prior probabilities

The observable state variable $o(t)$ comprises the data collected during the interval $t \leq t' \leq t + T_{\text{drift}}$. Formally it is the vector $[x(t'_0), \dots, x(t'_{N_{\text{drift}}})]$, whose dimension $N_{\text{drift}} + 1$ equals the interferometer sampling frequency ($\approx 16 \text{ kHz}$) times T_{drift} , where $x(t) = h(t) + n(t)$ is the output of the interferometer signal channel, $h(t)$ denotes the gravitational wave strain, and $n(t)$ denotes the noise. The emission probability matrix is then given by

$$L_{o(t),q_i} = \Pr[o(t)|f_{*i} \leq f_*(t) \leq f_{*i} + \Delta f_{\text{drift}}] \quad (8)$$

$$\propto \exp[G(f_{*i})], \quad (9)$$

where (9) follows from (8) by the definition of the frequency domain estimator.

Since we have no initial knowledge of $f_*(t)$, as prior we chose uniform distribution, i.e

$$\Pi_{q_i} = N_Q^{-1}, \quad (10)$$

for all q_i . The maximum entropy of uniform distribution serves us to avoid any unwarranted assumptions about the signal.

In binary neutron star applications, $L_{o(t),q_i}$ depends on a_0 as well as f_* ; an explicit formula is given in Section IV A. The prior is the product of a uniform prior in f_* and

² If we track amplitude $h(t)$ and frequency $f_*(t)$ jointly (outside the scope of this paper), then the state space enlarges appropriately, with $q(t) = [f_*(t), h(t)]$.

³ In one simple glitch model, let $\sigma = T_{\text{drift}}/T_{\text{glitch}}$ be the probability of a glitch occurring in an interval T_{drift} , where T_{glitch} is the mean glitch waiting time, and let η be the maximum fractional glitch size. Then we replace $1/3$ in equation (7) with $A_{q_{i-1}q_i} = (1 - \sigma)/3$, $A_{q_{i+1}q_i} = A_{q_iq_i} = (1 - \sigma)/3 + \sigma(\eta q_i / \Delta f_{\text{drift}})^{-1}$ and also have $A_{q_jq_i} = \sigma(\eta q_i / \Delta f_{\text{drift}})^{-1}$ for $1 < q_j - q_i \leq \eta q_i$.

a uniform or Gaussian prior in a_0 , the latter centred on the electromagnetically measured value with a standard deviation equal to the measurement uncertainty.

D. Viterbi algorithm

Let $Q^{*(k)} = [q^*(t_0), \dots, q^*(t_k)]$ be the first $k+1$ steps in the most probable path maximizing (4) for the observation sequence $O^{(k)} = [o(t_0), \dots, o(t_k)]$. For a Markov process, $Q^{*(k)}$ satisfies the nesting property $Q^{*(k-1)} \subset Q^{*(k)}$, i.e. $Q^{*(k-1)}$ makes up the first k steps of $Q^{*(k)}$. This is a special case of the Principle of Optimality [60]: if $Q^{*(k)}$ is optimal, then all subpaths within $Q^{*(k)}$ must be optimal too. In other words, given the optimal path $Q^{*(k)} = [q^*(t_0), \dots, q^*(t_i), \dots, q^*(t_j), \dots, q^*(t_k)]$, one must have $\Pr\{[q'(t_i), \dots, q'(t_j)] | [o(t_i), \dots, o(t_j)]\} \leq \Pr\{[q^*(t_i), \dots, q^*(t_j)] | [o(t_i), \dots, o(t_j)]\}$ for all possible choices of $[q'(t_i), \dots, q'(t_j)]$.

When applied to equation (4), the Principle of Optimality naturally defines a recursive algorithm, proposed by Viterbi [53], to find Q^* by backtracking. At every forward step $k > 0$ in the recursion, the Viterbi algorithm eliminates all but N_Q possible state sequences. The retained sequences end in different states by construction; if more than one sequence ends in a given state, the sequence with maximum $P[Q^{(k)} | O^{(k)}]$ is retained.

At time t_k , we save each of the N_Q maximum probabilities in the vector $\delta(t_k)$ with components

$$\delta_{q_i}(t_k) = \max_{q_j} \Pr[q(t_k) = q_i | q(t_{k-1}) = q_j; O^{(k)}] \quad (11)$$

and we save the state at t_{k-1} leading to each retained sequence in the vector $\Phi(t_k)$ with components

$$\Phi_{q_i}(t_k) = \arg \max_{q_j} \Pr[q(t_k) = q_i | q(t_{k-1}) = q_j; O^{(k)}], \quad (12)$$

with

$$\Pr[q(t_k) = q_i | q(t_{k-1}) = q_j; O^{(k)}] = L_{o(t_k)q_i} A_{q_i q_j} \delta_{q_j}(t_{k-1}) \quad (13)$$

in both (11) and (12). After stepping k forward through $1 \leq k \leq N_T$, we backtrack through the $\Phi(t_k)$ vectors to find the optimal path.

In summary, therefore, the Viterbi algorithm comprises four stages.

1. Initialization:

$$\delta_{q_i}(t_0) = L_{o(t_0)q_i} \Pi_{q_i}, \quad (14)$$

for $1 \leq i \leq N_Q$. Note that $\Phi_{q_i}(t_0)$ is never used.

2. Recursion:

$$\delta_{q_i}(t_k) = L_{o(t_k)q_i} \max_{1 \leq j \leq N_Q} [A_{q_i q_j} \delta_{q_j}(t_{k-1})], \quad (15)$$

$$\Phi_{q_i}(t_k) = \arg \max_{1 \leq j \leq N_Q} [A_{q_i q_j} \delta_{q_j}(t_{k-1})], \quad (16)$$

for $1 \leq i \leq N_Q$ and $1 \leq k \leq N_T$.

3. Termination:

$$\max P(Q|O) = \max_{q_j} \delta_{q_j}(t_{N_T}) \quad (17)$$

$$q^*(t_{N_T}) = \arg \max_{q_j} \delta_{q_j}(t_{N_T}) \quad (18)$$

for $1 \leq j \leq N_Q$.

4. Optimal path backtracking:

$$q^*(t_k) = \Phi_{q^*(t_{k+1})}(t_{k+1}) \quad (19)$$

for $0 \leq k \leq N_T - 1$.

By pruning the tree of possible paths efficiently at each step, the Viterbi algorithm reduces the number of comparisons from $N_Q^{N_T+1}$ to $(N_T + 1)N_Q^2$, which can be reduced further to $(N_T + 1)N_Q \ln N_Q$ by binary maximization [45].

III. ISOLATED NEUTRON STAR

We first consider an isolated neutron star, whose spin frequency wanders randomly by at most plus or minus one \mathcal{F} -statistic frequency bin on the time-scale T_{drift} . We take $T_{\text{obs}} = 370$ d, $T_{\text{drift}} = 10$ d, and $\Delta f_{\text{drift}} = 5.787037 \times 10^{-7}$ Hz for the sake of illustration and with an eye to later comparison with the tests in Section IV and V motivated by binaries like Sco X-1.

A. Matched filter: \mathcal{F} -statistic

Over time intervals that are short compared to T_{drift} , the optimal matched filter for a biaxial rotor with no orbital motion is the maximum-likelihood \mathcal{F} -statistic, which accounts for the rotation of the Earth and its orbit around the solar system barycentre (SSB).

The time-dependent data $x(t)$ collected at a single detector take the form

$$x(t) = \mathcal{A}^\mu h_\mu(t) + n(t), \quad (20)$$

where $n(t)$ represents stationary, additive noise, $h_\mu(t)$ are the four linearly independent signal components

$$h_1(t) = a(t) \cos \Phi(t), \quad (21)$$

$$h_2(t) = b(t) \cos \Phi(t), \quad (22)$$

$$h_3(t) = a(t) \sin \Phi(t), \quad (23)$$

$$h_4(t) = b(t) \sin \Phi(t), \quad (24)$$

$a(t)$ and $b(t)$ are the antenna-pattern functions defined by Equations (12) and (13) in Ref. [7], \mathcal{A}^μ is the amplitude associated with h_μ , and $\Phi(t)$ is the signal phase at the detector.

The \mathcal{F} -statistic is a frequency-domain estimator maximizing the likelihood of detecting a signal in noise with

respect to the four amplitudes \mathcal{A}^μ [7]⁴. It is defined as

$$\mathcal{F} = \frac{1}{2} x_\mu \mathcal{M}^{\mu\nu} x_\nu, \quad (25)$$

where we write $x_\mu = (x|h_\mu)$, and $\mathcal{M}^{\mu\nu}$ denotes the inverse matrix of $\mathcal{M}_{\mu\nu} = (h_\mu|h_\nu)$. The scalar product $(\cdot|\cdot)$ featuring in the definitions of x_μ and $\mathcal{M}_{\mu\nu}$ is defined as a sum over single-detector scalar products,

$$\begin{aligned} (x|y) &= \sum_X (x^X|y^X) \\ &= \sum_X 4\Re \int_0^\infty df \frac{\tilde{x}^X(f) \tilde{y}^{X*}(f)}{S^X(f)}, \end{aligned} \quad (26)$$

where X indexes the detector, $S^X(f)$ is the one-sided noise spectral density of detector X , the tilde denotes a Fourier transform, and \Re returns the real part of a complex number [61].

The expectation value of the \mathcal{F} -statistic is

$$E[2\mathcal{F}] = 4 + \rho_0^2, \quad (27)$$

where

$$\rho_0 = (\mathcal{A}^\mu \mathcal{M}_{\mu\nu} \mathcal{A}^\nu)^{1/2} \quad (28)$$

stands for the optimal SNR given a signal in Gaussian noise. The random variable $2\mathcal{F}$ is distributed according to a non-central chi-squared probability density function with four degrees of freedom, $p(2\mathcal{F}) = \chi^2(2\mathcal{F}; 4, \rho_0^2)$, whose non-centrality parameter ρ_0^2 is related to the amplitudes by

$$\rho_0^2 = \frac{1}{2} [A(\mathcal{A}_1^2 + \mathcal{A}_3^2) + B(\mathcal{A}_2^2 + \mathcal{A}_4^2) + 2C(\mathcal{A}_1\mathcal{A}_2 + \mathcal{A}_3\mathcal{A}_4)], \quad (29)$$

with $A = (a|a)$, $B = (b|b)$, and $C = (a|b)$. When there is no signal, the probability density function centralizes to give $p(2\mathcal{F}) = \chi^2(2\mathcal{F}; 4, 0)$. It can be shown that ρ_0 equals the SNR, with

$$\rho_0^2 = \frac{K h_0^2 T_{\text{drift}}}{S_n(2f_\star)}, \quad (30)$$

where the constant K depends on the right ascension, declination, polarization and inclination angles of the source. Averaging without bias over these angles yields $K = 4/25$ for a perpendicular rotor and an interferometer with perpendicular arms [7].

⁴ Equations (21)–(24) assume emission at $2f_\star$ only, i.e. $\Phi(t) \approx 4\pi f_\star t$, as appropriate for a perpendicular rotor. If the star's wobble angle is less than 90° , emission also occurs at f_\star , and eight amplitudes are involved, i.e. h_1, \dots, h_8 .

Parameter	Value	Units
f_\star	111.1	Hz
\dot{f}_\star	0.0	Hz s ⁻¹
ψ	4.08407	rad
$\cos \iota$	0.71934	–
ϕ_0	0.0	rad
α	4.27570	rad
δ	−0.27297	rad
$S_n(2f_\star)^{1/2}$	4×10^{-24}	Hz ^{-1/2}

TABLE I. Injection parameters used to create the synthetic data analysed in Sections III B and IV B.

B. Detectability versus h_0

We begin by illustrating the performance of the Viterbi algorithm with some representative examples. Seven sets of synthetic data are created for $T_{\text{obs}} = 370$ d at two detectors (H1 and L1) with $1 \leq h_0/(10^{-26}) \leq 20$ superposed on noise at a level typical of Advanced LIGO's design sensitivity, viz. $S_n(2f_\star)^{1/2} = 4 \times 10^{-24} \text{ Hz}^{-1/2}$, near the instrument's most sensitive frequency [54, 62]. Sky position and source orientation are specified in Table I. The synthetic SFTs are generated using *Makefake-data* version 4 from the LIGO data analysis software suite LALApps⁵. For each test, we take $T_{\text{drift}} = 10$ d, divide one year of data into $N_{T+1} = 37$ segments, and create a 1-Hz band of \mathcal{F} -statistic output containing the injected $f_\star(t)$ for each segment. We take $\Pi_{q_i} = N_Q^{-1} = \Delta f_{\text{drift}}$ as the prior and $\ln L_{o(t), q_i} = \mathcal{F}$ from the \mathcal{F} -statistic output for each 10-day segment. We also use the transition probability matrix given by equation (7). For each data set, we calculate the root-mean-square deviation ε (in Hz) between the optimal Viterbi path and the injected $f_\star(t)$. Tracking is deemed successful, if ε is smaller than one \mathcal{F} -statistic frequency bin (width Δf_{drift}). A systematic Monte-Carlo calculation of the threshold for detection lies outside the scope of this paper.

The outcomes of the tests above are presented in Table II and Figure 1. In Figures 1(a)–1(f), we see examples where the injected $f_\star(t)$ agrees closely with the optimal path reconstructed by the Viterbi algorithm. For $h_0/10^{-26} = 20, 10, 8, 6, 4, 2$, the maximum root-mean-square error is $0.47\Delta f_{\text{drift}}$. It arises mostly because the HMM takes one frequency bin as the smallest step, while the injected $f_\star(t)$ jumps to any value within ± 1 bin. For $h_0 = 1 \times 10^{-26}$, the optimal Viterbi path is not a good match; the maximum error is $\varepsilon = 2 \times 10^5 \Delta f_{\text{drift}}$. Indeed the closest match to the injected signal is the 700-th Viterbi path [see Figure 1(g)], and even then the match is worse than that for all the tests with $h_0 \geq 2 \times 10^{-26}$ ($\varepsilon = 2.15\Delta f_{\text{drift}}$).

The rapid loss of detectability experienced at $h_0 \approx 1 \times 10^{-26}$ in Table II is expected theoretically. Figure 2

⁵ <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>

$h_0(10^{-26})$	Detect?	ε (Hz)	$\varepsilon/\Delta f_{\text{drift}}$
20.0	✓	1.866×10^{-7}	0.322
10.0	✓	1.663×10^{-7}	0.287
8.0	✓	2.132×10^{-7}	0.368
6.0	✓	1.683×10^{-7}	0.291
4.0	✓	2.054×10^{-7}	0.355
2.0	✓	2.719×10^{-7}	0.470
1.0	×	0.112	2×10^5

TABLE II. Outcome of Viterbi tracking for injected signals from isolated sources with the parameters in Table I, $T_{\text{obs}} = 370$ d, $T_{\text{drift}} = 10$ d, and wave strain h_0 . The root-mean-square error ε between the optimal Viterbi track and injected $f_*(t)$ is quoted in Hz and in units of Δf_{drift} , the \mathcal{F} -statistic frequency bin width.

displays the formal root-mean-square error computed numerically as a function of N_T for fixed ρ_0^2 [Figure 2(a)] and as a function of ρ_0^2 for fixed N_T [Figure 2(b)]. The error is approximated as a linear combination of two terms, one arising from the probability of an outlier (see Sections III C and III D), and the other set by the Cramér-Rao (CR) lower bound [63]. For large SNR, the error variance of the estimator approaches the CR bound, i.e. the nearly horizontal, rightmost segments of the blue and green curves. As the SNR decreases, the error variance departs from the CR bound more and more and eventually becomes unbounded. For example, looking at the red curve in Figure 2(a) [for $\rho_0^2 = 1$, corresponding to the injection in Figure 1(g) with $h_0 = 1 \times 10^{-26}$], we see that the error variance does not asymptotically approach the CR bound over the plotted range of N_T . By contrast, in Figure 2(b) for example, the green curve (for $N_T = 36$) hugs the CR bound for $\rho_0^2 \gtrsim 6$, corresponding to $h_0 > 2.6 \times 10^{-26}$. However, as the SNR decreases in the regime $\rho_0^2 \lesssim 6$, the error diverges rapidly away from the turning point in the green curve, with $\varepsilon \approx 10^{-0.5\rho_0^3}$ Hz as a rough approximation. This behaviour is typical at low SNR near the detection boundary [63]. The probability that the optimal Viterbi path coincides with the injection in Figure 1(g) is approximately 35%, if the test is repeated for a large number of realizations of the noise. We quantify the probability that the optimal Viterbi path matches the injection in Section III D. Notice that, as always, extending T_{obs} or increasing T_{drift} (if the wandering is slow enough) improves the tracking.

C. Distribution of path probabilities

The logarithm of equation (4) expresses $\ln P(Q|O)$ as a sum of log likelihoods, each of which is chi-squared-distributed with four degrees of freedom, if $G(f_*)$ is the \mathcal{F} -statistic. As the chi-squared distribution is additive, we can easily calculate the probability density function of $z = P[Q^{(k)}|O^{(k)}]$ as a function of step number k , with $0 \leq k \leq N_T$, along any Viterbi path. If the Viterbi path coincides exactly with the true path, we obtain the

probability density function

$$p(z) = \chi^2(z; 4k, k\rho_0^2). \quad (31)$$

If the Viterbi path does not intersect the true path anywhere, we obtain

$$p(z) = \chi^2(z; 4k, 0). \quad (32)$$

If the Viterbi path intersects the true path at some steps but not others, $p(z)$ lies somewhere between $\chi^2(z; 4k, k\rho_0^2)$ and $\chi^2(z; 4k, 0)$. Note that $p(z)$ for the optimal Viterbi path lies somewhere between the above bounds but its functional form differs from (31) and (32). The Viterbi algorithm maximizes z over all 3^k paths of length k terminating at a given frequency bin $q(t_k)$. Hence, for pure noise, $p(z)$ for the optimal Viterbi path is constructed from (32) via the extreme value theorem modified to account for the fact that the paths overlap and are therefore correlated. This calculation is hard to do analytically and is postponed to future work.

Figure 3 displays $p(z) = \chi^2(z; 4k, k\rho_0^2)$ (perfect intersection) and $\chi^2(z; 4k, 0)$ (no intersection) for the representative example $h_0 = 2 \times 10^{-26}$ (i.e. $\rho_0^2 = 3.6$) in Table II. The graph demonstrates clearly how it is progressively easier to detect a signal, as more Viterbi steps are taken. After $k = 1$ steps, $\chi^2(z; 4k, k\rho_0^2)$ and $\chi^2(z; 4k, 0)$ are hard to distinguish, with the difference confined to the tail. After $k = 10$ steps, $\chi^2(z; 4k, k\rho_0^2)$ and $\chi^2(z; 4k, 0)$ are well separated everywhere, including at the peak.

D. Probability of no outlier

In order for the optimal Viterbi path to be a reliable detection agent, we desire a high probability $1 - P_{\text{outlier}}$ that the frequency bin containing the signal returns a higher value of $G(f)$ than all the other $N_Q - 1$ frequency bins, which do not contain a signal, i.e. there are no outliers. Mathematically this translates to the condition $\delta_{q_i}(t_{N_T}) < \delta_{q_s}(t_{N_T})$ for all $i \neq s$ and for all possible values of the measurement $\delta_{q_s}(t_{N_T})$ in the state q_s that contains the signal. From Section III C and the Optimality Principle we obtain

$$1 - P_{\text{outlier}} = \int_0^\infty dz' \chi^2(z'; 4N_T, N_T\rho_0^2) P(z'; 4N_T, 0)^{N_Q-1}, \quad (33)$$

where $P(z; \alpha, \beta) = \int_0^z dz' \chi^2(z'; \alpha, \beta)$ is the chi-squared cumulative distribution function.

Figure 4 shows the probability of no outlier computed numerically as a function of N_T for fixed ρ_0^2 [Figure 4(a)] and as a function of ρ_0^2 for fixed N_T [Figure 4(b)]. Detectability improves with the number of Viterbi steps. From the shape of the curves, it is clear that the lower the value of ρ_0^2 the more steps N_T are required.

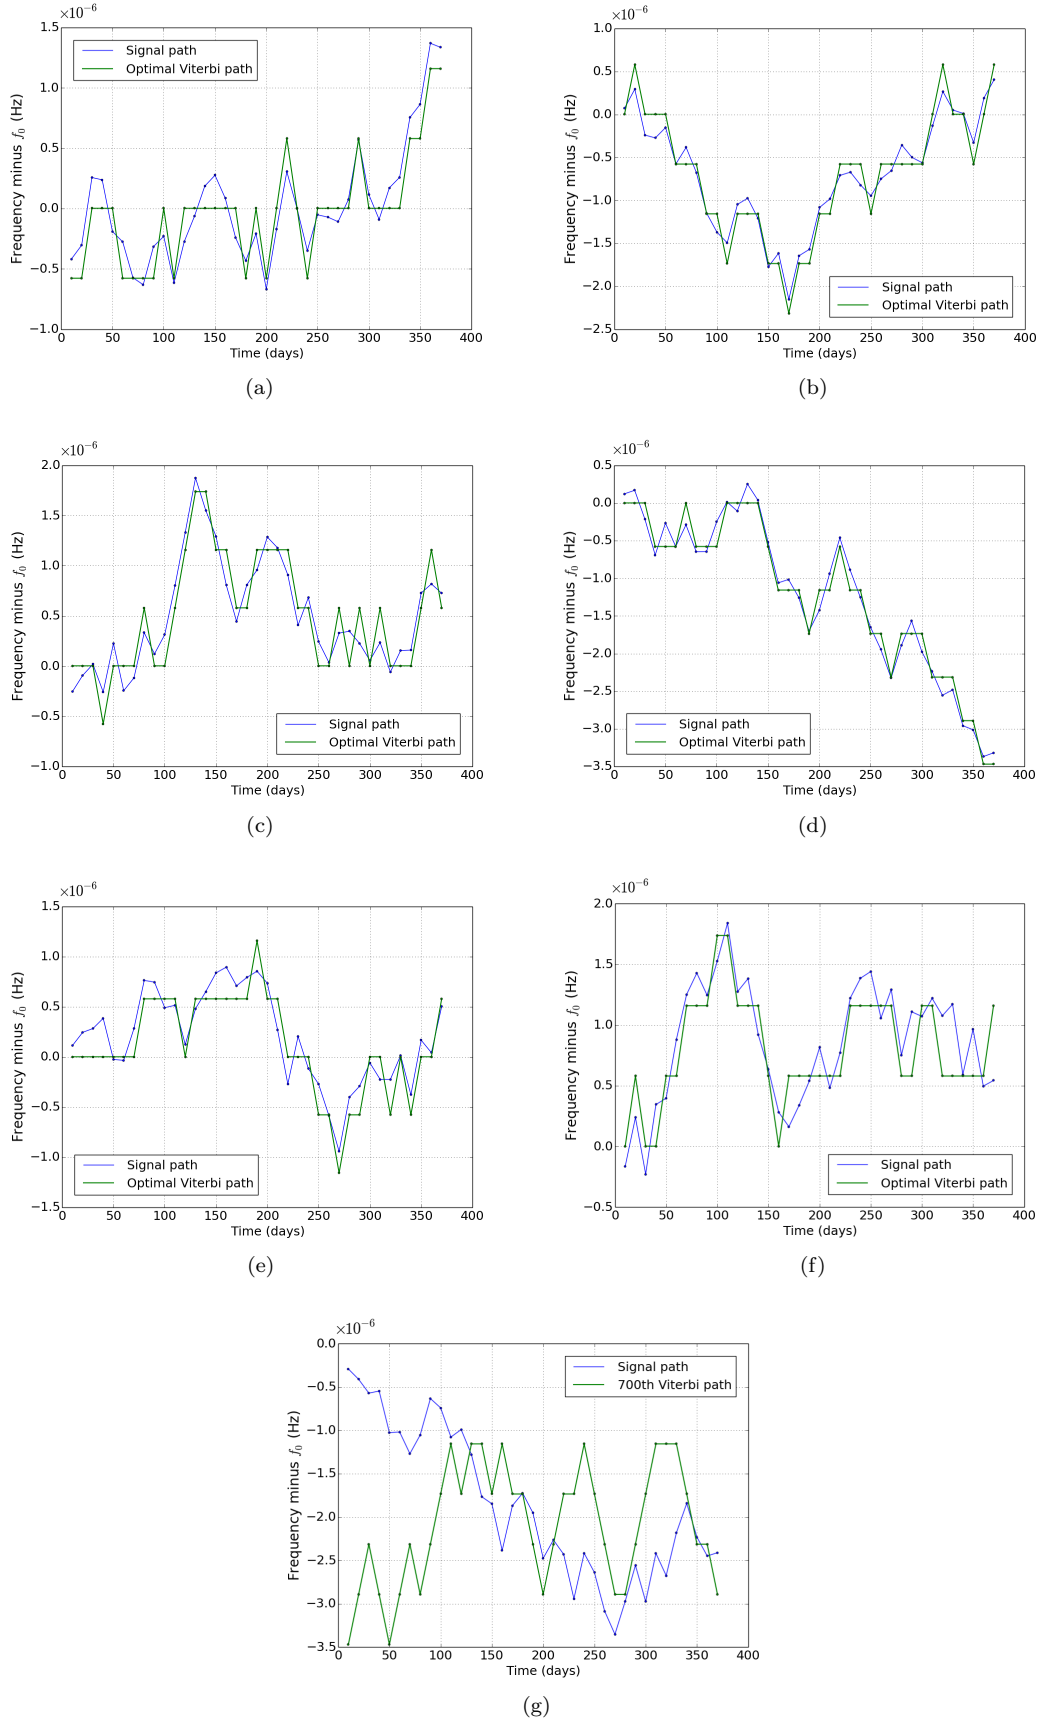


FIG. 1. True $f_\star(t)$ (blue curve) and Viterbi path (green curve) for the seven injected signals in Table II, representing an isolated neutron star. Panels (a)–(f) display optimal Viterbi paths for $h_0/10^{-26} = 20, 10, 8, 6, 4, 2$ respectively; a good match is obtained in each case. Panel (g) displays the closest-matching path (i.e. smallest ϵ) for $h_0 = 1 \times 10^{-26}$; the match is poor. The units on the horizontal (time) and vertical (frequency) axes are days and Hz respectively. The symbol f_0 in the vertical axis label stands for $f_\star(0)$.




FIG. 2. Formal root-mean-square error estimate in the vicinity of the detection threshold as a function of (a) N_T , with fixed $\rho_0^2 = 10, 7$ and 1 , corresponding to $h_0 = 3.4 \times 10^{-26}, 2.8 \times 10^{-26}$ and 1×10^{-26} , and (b) ρ_0^2 , with fixed $N_T = 50, 36$ and 1 . The nearly horizontal, rightmost parts of the blue and green curves correspond to the CR lower bound [63].




FIG. 3. Probability density function $p(z)$ for $z = P[Q^{(k)}|O^{(k)}]$ after $k = 1$ (red curve) and $k = 10$ (blue curve) Viterbi steps, assuming $Q^{(k)}$ intersects either perfectly (solid curve) or not at all (dashed curve) with the true path $f_*(t)$. Parameters: $\rho_0^2 = 3.6$; corresponding to $h_0 = 2 \times 10^{-26}$ in Table II.

IV. BINARY NEUTRON STAR

A. Matched filter: Bessel-weighted \mathcal{F} -statistic

When a biaxial rotor orbits a binary companion, the gravitational wave strain is frequency modulated due to the orbital Doppler effect. The signal at the detector is given by

$$x(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t) + n(t), \quad (34)$$

where F_+ and F_\times are the beam-pattern functions defined in Equations (10) and (11) in Ref. [7]. For a Keplerian

orbit, one has

$$h_{+, \times}(t) \propto \cos\{2\pi f_*[t + a_0 \sin(2\pi t/P)]\}, \quad (35)$$

where a_0 is the projected semimajor axis, and P is the orbital period. Expanding equation (35) by the Jacobi-Anger identity [64], we obtain

$$h_{+, \times}(t) \propto \sum_{n=-\infty}^{\infty} J_n(2\pi f_* a_0) \cos[2\pi(f_* + n/P)t], \quad (36)$$

where $J_n(z)$ is a Bessel function of order n of the first kind.

The coefficients $J_n(z)$ in equation (36) decay rapidly for $|n| > z$ ($z = 2\pi f_* a_0 \gg 1$), and the gravitational wave power is distributed into approximately $M = 2\text{ceil}(2\pi f_* a_0) + 1$ orbital sidebands separated by $1/P$, where $\text{ceil}(x)$ denotes the smallest integer greater than or equal to x . Equation (36) suggests that, over time intervals that are short compared to T_{drift} , the optimal matched filter takes the form of a convolution

$$G(f) = \mathcal{F}(f) \otimes B(f), \quad (37)$$

where $B(f)$ is the squared modulus of the Fourier transform of the sum in equation (36) heterodyned at f_* , viz.

$$B(f) = \sum_{n=-(M-1)/2}^{(M-1)/2} [J_n(2\pi f a_0)]^2 \delta(f - n/P). \quad (38)$$

Let us now estimate approximately how the SNR depends on ρ_0 , f_* , and a_0 . For the purpose of the following calculation we write $\mathcal{F}(f) \approx \rho_0^2 B(f) + W(f)$, where $W(f)$ is the \mathcal{F} -statistic of the noise, modeled as a chi-squared-distributed random variable with four degrees of freedom.

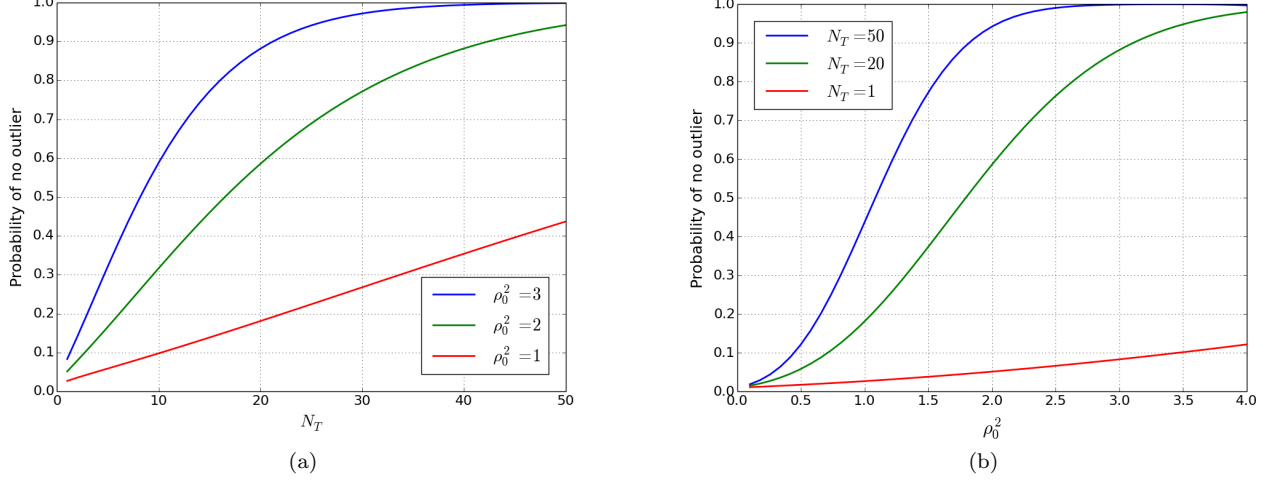


FIG. 4. Probability of no outlier computed numerically with $N_Q = 100$ \mathcal{F} -statistic frequency bins as a function of (a) N_T , with ρ_0^2 fixed, and (b) ρ_0^2 , with N_T fixed. Note that we have $N_Q \sim 10^6$ in real applications, but the integral in equation (33) is hard to compute numerically in this regime. Fortunately the curve depends weakly on N_Q for $N_Q \gg 1$, with $1 - P_{\text{outlier}}$ varying by $\sim 4\%$ from $N_Q = 100$ to $N_Q = 200$.

This implies in particular $\text{var}(W) = 8$, where var denotes the variance. The SNR yielded by $G(f)$ evaluated at the source frequency $f = f_*$ is given by

$$\text{SNR} = \frac{\rho_0^2 (B \otimes B)(f_*)}{[\text{var}(W \otimes B)]^{1/2}}. \quad (39)$$

For large M , $W \otimes B$ is approximately Gaussian, with variance

$$\text{var}(W \otimes B) \approx 8 \sum_{n=-\infty}^{\infty} J_n^4(2\pi f_* a_0), \quad (40)$$

and we also have

$$(B \otimes B)(f_*) = \sum_{n=-\infty}^{\infty} J_n^4(2\pi f_* a_0), \quad (41)$$

implying

$$\text{SNR} = \frac{\rho_0^2}{2\sqrt{2}} \left[\sum_{n=-\infty}^{\infty} J_n^4(2\pi f_* a_0) \right]^{1/2}. \quad (42)$$

As $\sum_{n=-\infty}^{\infty} J_n^4(z)$ is bounded by z^{-1} for large z [64], we infer the lower bound

$$\text{SNR} \geq \frac{\rho_0^2}{(16\pi f_* a_0)^{1/2}}. \quad (43)$$

Note, this inequality requires $f_* > 40$ Hz, which is adequate for our purpose. In previous frequency domain searches for binaries, e.g. Sco X-1 [55], the matched filter (37) and (38) is replaced by an unweighted comb of

orbital sidebands of the form

$$B_{\text{comb}}(f) = \frac{1}{M} \sum_{n=-(M-1)/2}^{(M-1)/2} \delta(f - n/P), \quad (44)$$

called the \mathcal{C} -statistic. By an argument similar to the one in the previous paragraph, we have

$$\rho_0^2 (B \otimes B_{\text{comb}})(f_*) = \rho_0^2, \quad (45)$$

$$\text{var}(W \otimes B_{\text{comb}}) = 32\pi f_* a_0, \quad (46)$$

and hence

$$\text{SNR} \geq \frac{\rho_0^2}{(32\pi f_* a_0)^{1/2}}. \quad (47)$$

Equations (43) and (47) demonstrate that the Bessel-weighted matched filter can recover approximately $\sqrt{2}$ times the SNR of the \mathcal{C} -statistic. Unlike for an isolated source, where the SNR depends only on ρ_0 , the SNR for a binary source is also inversely proportional to $\sqrt{f_*}$ (i.e. $\text{SNR} \propto h_0/\sqrt{f_*}$), adversely affecting performance at higher frequencies.

Figure 5 shows an example comparing the performance of the orbital sideband filters in equation (44) (left panels; unweighted) and (38) (right panels; Bessel weighted) on a 10-day data segment with an injected signal at $f_*(t) = 731.0068$ Hz from a binary source. The Bessel-weighted filter takes $P = 68023.7$ s and $a \sin i = 1.44$, i.e. values characteristic of Sco X-1 (see Sections IV B and V). Panels (a) and (b) (top and bottom) show 1-Hz and 0.015-Hz frequency bands containing the signal respectively. The injected frequency is marked by a red, vertical

dashed line. Not only does the Bessel-weighted filter recover more signal power, but also the structure of its peak suits Viterbi tracking better. When we zoom into the peak [Figure 5(b); 731.0–731.015 Hz], the \mathcal{C} -statistic output is relatively flat over a band of width $\approx 2M\Delta f_{\text{drift}}$, presenting the Viterbi tracker with multiple options, each with relatively low SNR. By contrast the Bessel-weighted filter marshals more of the power into a single, distinct peak, which is easier for the Viterbi algorithm to track.

B. Detectability versus h_0

We begin by illustrating the performance of the Viterbi tracker for binary sources with some representative examples. We inject signals into Gaussian noise and generate synthetic SFTs for $T_{\text{obs}} = 370$ d at two interferometers using *Makefakedata* version 4 as described in Section III. We keep the same source parameters as in Section III B and introduce the orbital parameters listed in Table III, copied from Sco X-1 for definiteness. The analysis for each realisation proceeds in three steps. (1) We calculate the \mathcal{F} -statistic in a 1-Hz band containing the injection for each segment ($T_{\text{drift}} = 10$ d) of data and output 37 segments for a year. (2) We create a Bessel-weighted filter [equation (38)] with $P = 68023.7$ s and $a \sin i = 1.44$ s, process each 10-day \mathcal{F} -statistic segment, and generate 37 $G(f)$ outputs. (3) We apply the Viterbi tracker to the $G(f)$ output and find the optimal Viterbi path like in Section III B.

In realistic applications, a_0 is often known approximately but not exactly from electromagnetic observations (typical uncertainty $\approx \pm 10\%$) [54, 65]. Hence in general we track a_0 as well as $f_*(t)$, as described in Sections II B and II C. However, for the tests in this section, our initial guess for a_0 matches exactly the injected value in Table III, so we only track $f_*(t)$. (In Section V, a_0 is tracked too.)

The Bessel-weighted filter in equation (38) depends on $2\pi f a_0$, complicating step (2) in the paragraphs above. Strictly speaking, the filter takes a slightly different form in each \mathcal{F} -statistic frequency bin within each 1-Hz band, which would be prohibitive computationally to implement. Instead, we execute step (2) using a filter with Bessel weightings $J_n(2\pi \bar{f} a_0)$, where \bar{f} is the central frequency in each 1-Hz band. The fractional error thereby introduced across the 1-Hz band is minimal ($\lesssim 1\%$) compared to the fractional uncertainty in a_0 from electromagnetic observations. A similar approach was adopted in previous \mathcal{C} -statistic searches [54–56]. The fractional uncertainty in P from electromagnetic observations is typically a few parts in 10^7 and can be neglected.

The test outcomes are presented in Table IV and Figure 6. Figures 6(a)–6(d) show the tracking results for $h_0/10^{-26} = 20, 10, 8, 6$. Tracking is deemed successful, if the root-mean-square discrepancy ε between the optimal Viterbi path and injected $f_*(t)$ is less than one $G(f)$ frequency bin (width Δf_{drift}) or the width of the

orbital sideband pattern (M/P), whichever is larger. In Figure 6(a)–6(c), the injected $f_*(t)$ agrees well with the optimal Viterbi path. For $h_0/10^{-26} = 20, 10, 8$, the maximum root-mean-square error is $0.814 \Delta f_{\text{drift}}$. For $h_0 = 6 \times 10^{-26}$, the optimal Viterbi path is a poor match with $\varepsilon = 7 \times 10^5 \Delta f_{\text{drift}}$. The closest match to the injected signal is the seventh Viterbi path with $\varepsilon = 7.677 \times 10^{-7} \text{ Hz} = 1.327 \Delta f_{\text{drift}}$ [see Figure 6(d)]. In other words, the sensitivity drops four-fold from $h_0 \approx 2 \times 10^{-26}$ for an isolated source to $h_0 \approx 8 \times 10^{-26}$ for a binary source.

We quantify the error in tracking f_* as a function of the error in the assumed value of a_0 , by injecting a strong signal with $h_0/\sqrt{S_n(2f_*)} = 10^7 \text{ Hz}^{1/2}$ into a 10-day segment, and tracking f_* as well as a_0 over 250 bins spanning $\pm 25\%$ of the injected value a_0^{true} . The results are displayed in Figure 7, a contour plot displaying the log likelihood $\ln P[Q^*(O)|O]$ as a function of a_0 (expressed as the percentage offset from the true, injected value a_0^{true}) and f_* (expressed as the absolute offset from the true, injected value f_*^{true}). The bright colors (cyan, yellow, red) stand for the highest log likelihoods. The maximum value (in red) is found at the injected values of f_*^{true} and a_0^{true} . Caused by the $\pm 25\%$ uncertainty in a_0 , the maximum absolute offset of f_* recovered by the tracker from f_*^{true} is $\varepsilon_{f_*} \approx 0.004 \text{ Hz}$ ($0.135 M/P$). In reality the $\pm 10\%$ uncertainty in a_0 leads to an uncertainty in f_* , given by $\varepsilon_{f_*} \approx 0.001 \text{ Hz}$ ($0.003 M/P$).

V. SCO X-1 MOCK DATA CHALLENGE: A “REALISTIC” EXAMPLE

In this section, we combine the Viterbi tracker and Bessel-weighted matched filter validated in Section IV to search for the 50 Sco-X-1-type signals (50–1500 Hz) generated for the Sco X-1 Mock Data Challenge (Version 6) [54]. The aims of the exercise are three-fold: (i) to test the performance of the Viterbi tracker under “realistic” conditions on a data set generated by an independent party; (ii) to compare the performance of the Viterbi tracker against the CrossCorr [9, 10, 66], TwoSpect [21], Radiometer [67–69], Sideband [55, 70] and Polynomial [71] pipelines which competed in the Mock Data Challenge; and (iii) to prepare the Viterbi tracker for Advanced LIGO observations.

The parameters of the 50 injected signals in the Mock Data Challenge are listed in Table III in Ref. [54]. The 50 signals were originally “closed”, i.e. their parameters were kept secret, in order to compare blindly the competing pipelines from the perspectives of sensitivity, parameter estimation and efficiency. Four pipelines (TwoSpect, Radiometer, Sideband and Polynomial) competed under closed conditions in Ref. [54]. CrossCorr analysed the data in self-blinded mode, after the injection parameters were revealed. We note that $f_*(t)$ does not wander for any of the injected signals, a situation which the Viterbi tracker with the tridiagonal transition matrix in Equa-




FIG. 5. Convolution of the \mathcal{F} -statistic with a matched filter comprising unweighted (left panels) and Bessel-weighted (right panels) orbital sidebands for a 10-day data segment injected with a signal from a binary source. The red dashed lines indicate the injected frequency. The top and bottom panels show 1-Hz and 0.015-Hz bands around the injected frequency $f_*(t)$ respectively.

Parameter	Value	Units	Description
P	68023.7	s	Orbital period
a_0	1.44	s	Projected orbital semimajor axis
Δa_0	0.18	s	Measurement error in a_0
T_P	1245984672	s	Time of periaapsis passage in SSB
e	0.0	—	Orbital eccentricity

TABLE III. Orbital parameters used to create the synthetic data for the binary sources analysed in Section IV B. $2\Delta a_0$ is the width of the prior distribution of a_0 .

$h_0 (10^{-26})$	Detect?	ε (Hz)	$\varepsilon/\Delta f_{\text{drift}}$	$\varepsilon P/M$
20.0	✓	3.296×10^{-7}	0.570	1.114×10^{-5}
10.0	✓	4.655×10^{-7}	0.804	1.573×10^{-5}
8.0	✓	4.709×10^{-7}	0.814	1.591×10^{-5}
6.0	×	0.378	7×10^5	12.773

TABLE IV. Outcome of Viterbi tracking for injected signals from binary sources with the parameters in Tables I and III, $T_{\text{obs}} = 370$ d, $T_{\text{drift}} = 10$ d, and wave strain h_0 . The root-mean-square error ε between the optimal Viterbi track and injected $f_*(t)$ is quoted in Hz, in units of Δf_{drift} , the \mathcal{F} -statistic frequency bin width, and in units of M/P , the half-width of the orbital sideband pattern.

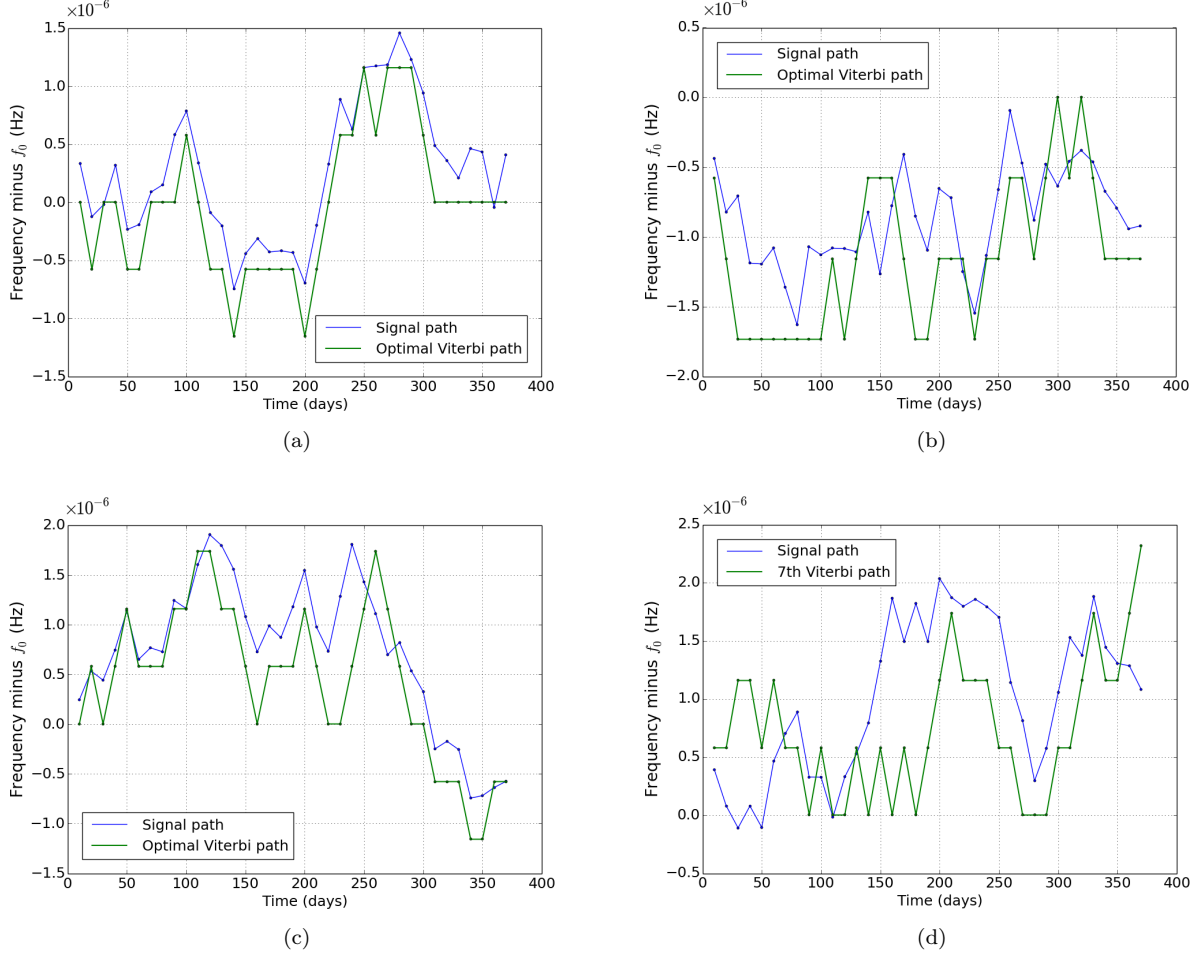


FIG. 6. True $f_\star(t)$ (blue curve) and Viterbi path (green curve) for the four injected signals in Table IV, representing a neutron star in a binary orbit. Panels (a)–(c) display optimal Viterbi paths for $h_0/10^{-26} = 20, 10, 8$ respectively; a good match is obtained in each case. Panel (d) is the closest-matching path (i.e. smallest ε) for $h_0 = 6 \times 10^{-26}$. The units on the horizontal (time) and vertical (frequency) axes are days and Hz respectively. The symbol f_0 in the vertical axis label stands for $f_\star(0)$.

tion (7) handles easily and without bias. To mimic a real search, we claim a detection if the log likelihood of the optimal Viterbi path, $\ln P[Q^*(O)|O]$, exceeds its mean value plus seven standard deviations. The choice of seven standard deviations in this paper is arbitrary but it is broadly consistent with the thresholds chosen in the previous Sideband searches for Sco X-1 in LIGO S5 data [72] and in Stage I of the MDC [54], yielding approximately the same detection rate using a 10-day segment. A more systematic Monte-Carlo calculation of the threshold and false alarm rate lies outside the scope of this paper.

We conduct the search in three stages. Firstly, we pick the same 10-day segment of MDC data analysed in Ref. [54] by the Sideband pipeline, starting at GPS time 1245000000. We find that 12 out of 50 signals are detected, matching the performance of the \mathcal{C} -statistic. Secondly, for the 38 signals that are not detected in 10 days, we analyse a one-year stretch of data starting at GPS 1230338490. We find that 23 extra signals are detected,

leaving 15 out of 50 undetected. The first two stages are performed with data from two interferometers due to computational limitations. In the third stage, we reanalyse the 15 remaining signals using three interferometers. Gratingly, we find that we detect six extra signals. Stages two and three are performed for $T_{\text{obs}} = 1 \text{ yr}$, i.e. on the same footing as the four non-Sideband algorithms competing in Ref. [54]. We present the results from the three stages in detail in Section V A, V B and V C below and tabulate them in Table VI.

We find that the error in the estimates of $f_\star(0)$ and a_0 , denoted by $\varepsilon_{f_\star(0)}$ and ε_{a_0} respectively, satisfy $\varepsilon_{f_\star(0)} < 7.4 \times 10^3 \Delta f_{\text{drift}}$ (i.e. $4.26 \times 10^{-3} \text{ Hz}$) and $\varepsilon_{a_0} < 0.4a_0$ in all cases where there is a successful detection.




FIG. 7. Contour plot of the log likelihood $\ln P[Q^*(O)|O]$ as a function of a_0 (expressed as the percentage offset from the true, injected value a_0^{true}) and f_* (expressed as the absolute offset from the true, injected value f_*^{true}) for a strong ($\text{SNR} \gg 1$) binary-star signal with constant f_* observed during a single 10-d segment ($T_{\text{obs}} = 10$ d). The bright colors (cyan, yellow, red) stand for the highest log likelihoods. The injected signal parameters are $f_*^{\text{true}} = 111.1$ Hz, $a_0^{\text{true}} = 1.44$, and $\Delta a_0 = 0.25a_0^{\text{true}}$, characteristic of Scorpius X-1.

A. $T_{\text{obs}} = 10$ d, two interferometers

In the first stage, we pick the same 10-day segment of MDC data analysed in Ref. [54] by the sideband pipeline from two interferometers (H1 and L1), starting at GPS time 1245000000. We search a 1-Hz frequency band containing the signal for each injection, setting $P = 68023.7$ s and tracking a 0.72-s band of a_0 centred on the electromagnetic observation value 1.44 s. A uniform prior is set for both f_* and a_0 . The first stage successfully detects injections 1, 3, 15, 20, 32, 35, 59, 62, 65, 66, 75 and 84. We detect 12 signals rather than the 16 found by the Sideband pipeline in Ref. [54], which used data from three interferometers. As a cross-check, we perform a supplementary search for the four missing signals with three interferometers and $T_{\text{obs}} = 10$ d and detect them all.

B. $T_{\text{obs}} = 1$ yr, two interferometers

For the 38 out of 50 signals that are not detected in a single 10-day segment, we do Viterbi tracking for $T_{\text{obs}} = 1$ yr using data from two interferometers (H1 and L1). The search space and prior are the same as those in Section V A. In this stage 23 extra injections are successfully detected: 2, 5, 11, 14, 17, 19, 23, 26, 29, 36, 44, 47, 51, 60, 61, 67, 68, 76, 79, 83, 85, 95, and 98.

C. $T_{\text{obs}} = 1$ yr, three interferometers

For the remaining 15 signals that are not detected in the first two stages, we do Viterbi tracking for $T_{\text{obs}} = 1$ yr using data from three interferometers (H1, L1 and V1). The search space and prior are the same as those in Sections V A and V B. In this last stage six out of the remaining 15 injections are successfully detected: 21, 50, 52, 54, 58, and 71.

Figure 8 shows the error in estimated f_* as a function of $h_0/\sqrt{f_*}$ for the 50 injected signals in Stage I of the Sco X-1 MDC. The circles, stars and triangles mark injections detected in stages one ($T_{\text{obs}} = 10$ d, two interferometers), two ($T_{\text{obs}} = 1$ yr, two interferometers) and three ($T_{\text{obs}} = 1$ yr, three interferometers) respectively. The squares mark the injections not detected in any of the three stages. All the nine undetected signals have low SNR, with $h_0/\sqrt{f_*} \lesssim 1 \times 10^{-26} \text{ Hz}^{-1/2}$. None are detected by any competing pipeline except for CrossCorr in Ref. [54]. Viterbi tracking detects seven more signals (21, 50, 52, 54, 58, 71, 98) than TwoSpect, with $4.7 \times 10^{-27} \text{ Hz}^{-1/2} \leq h_0/\sqrt{f_*} \leq 1.46 \times 10^{-26} \text{ Hz}^{-1/2}$ for these seven.

We have also verified that injections 41, 48, 57, 63, 64, 69, 72, 73 and 90 are not detected by the HMM, even when the nonwandering character of the MDC Stage I signals is recognized explicitly by choosing a diagonal transition matrix $A_{q_i q_j} = \delta_{q_i q_j}$. This confirms that the Viterbi algorithm finds wandering and nonwandering signals with approximately equal efficiency, as long as the number of possible transitions at each step is relatively small (three or less here).

VI. CONCLUSION

In this paper, we describe an HMM method for tracking a continuous gravitational wave signal with wandering spin frequency, emitted by either an isolated neutron star or a neutron star in a binary orbit. The HMM assumes a simple, nearest-neighbour-bin transition matrix combined with emission probabilities given by standard maximum likelihood matched filters: \mathcal{F} -statistic for an isolated target, and its Bessel-weighted version for a binary. The HMM is solved recursively for the optimal frequency history $f_*(t)$ using the Viterbi algorithm. It is shown that, for Gaussian noise at a level characteristic of Advanced LIGO and with total observation time $T_{\text{obs}} = 1$ yr, the algorithm successfully tracks signals with $h_0 \gtrsim 2 \times 10^{-26}$ (isolated), and $h_0 \gtrsim 8 \times 10^{-26}$ (binary).

When applied to Stage I of the Scorpius X-1 Mock Data Challenge, the Viterbi tracker successfully detects 41 out of 50 synthetic signals with $\varepsilon_{f_*(0)} < 4.26 \times 10^{-3}$ Hz. In comparison, the CrossCorr, TwoSpect, Radiometer, Sideband and Polynomial algorithms detected 50, 34, 28, 16 and 7 out of 50 signals respectively. Performance metrics are summarized in Table V. The frequency estimation error achieved by the Viterbi algorithm ranges from

	CrossCorr	Viterbi	TwoSpect	Radiometer	Sideband	Polynomial
Hit rate (out of 50)	50	41	34	28	16	7
Best h_0 (10^{-25})	0.684	1.093	1.250	2.237	3.565	7.678
Best $h_0/\sqrt{f_\star}$ ($\times 10^{-25}$ Hz $^{-1/2}$)	0.020	0.047	0.082	0.102	0.235	0.261
Typical ε_{f_\star} (Hz)	10^{-5}	$10^{-7} - 10^{-3}$	10^{-4}	10^{-1}	10^{-2}	10^{-2}
Typical ε_{a_0} (s)	10^{-4}	$10^{-6} - 10^{-1}$	10^{-2}	—	—	—
Typical run time (CPU-hr)	10^6	10^3	10^5	10^3	10^3	10^8

TABLE V. Comparison of Viterbi tracker and competing algorithms in Stage I of the Scorpius X-1 Mock Data Challenge [54].




FIG. 8. Performance of the Viterbi tracker in Stage I of the Sco X-1 MDC. Error in estimated f_\star as a function of $h_0/\sqrt{f_\star}$ for the 50 injected signals in Stage I of the Sco X-1 MDC. The blue circles, green stars and red triangles mark injections detected in stages one ($T_{\text{obs}} = 10$ d, two interferometers), two ($T_{\text{obs}} = 1$ yr, two interferometers) and three ($T_{\text{obs}} = 1$ yr, three interferometers) respectively. The yellow squares mark the injections not detected in any of the three stages. Injection parameters are provided in Table VI.

$\sim 10^{-3}$ Hz to $\sim 10^{-7}$ Hz in the first and second stages ($T_{\text{obs}} = 10$ d and 1 yr, two interferometers) and from $\sim 10^{-5}$ Hz to $\sim 10^{-7}$ Hz in the third stage ($T_{\text{obs}} = 1$ yr, three interferometers). In comparison, the frequency estimation errors achieved by the CrossCorr, TwoSpect, Radiometer, Sideband and Polynomial algorithms are of order 10^{-5} Hz, 10^{-4} Hz, 10^{-1} Hz, 10^{-2} Hz and 10^{-2} Hz respectively. The a_0 estimation error achieved by the Viterbi algorithm spans the range $10^{-6} \leq \varepsilon_{a_0}/(1 \text{ s}) \leq 0.6$. In comparison, only CrossCorr and TwoSpect estimate a_0 in the tests contained in Ref. [54], achieving $\varepsilon_{a_0} \sim 10^{-4}$ s and 10^{-2} s respectively.

One advantage of the Viterbi tracker with respect to its competitors in the Mock Data Challenge is computational speed. For a 1-Hz band and $T_{\text{obs}} = 1$ yr, it takes ~ 1 CPU-hr to create \mathcal{F} -statistic data for $N_T = 36$ 10-

d segments. It then takes ~ 0.3 CPU-hr to process the \mathcal{F} -statistic data with the Viterbi tracker. Tracking both f_\star and a_0 takes slightly longer than tracking f_\star only, depending on the number of a_0 bins. In contrast, the CrossCorr and TwoSpect algorithms, which detected 50 and 34 out of 50 synthetic signals respectively, require $\sim 10^6$ CPU-hr to complete a typical broadband (0.5-kHz) search [54]. The computational savings offered by the Viterbi tracker can be re-invested to extend the astrophysical goals of the search, e.g. by searching a larger parameter space for Scorpius X-1 or targeting other low mass X-ray binaries [73].

Stage I of the Mock Data Challenge did not involve spin wandering. Nonetheless Viterbi tracking works exactly the same way, whether $f_\star(t)$ wanders or not. The algorithm is blind to the exact form of the transition matrix, so there is every reason to expect that the minimum h_0 detectable in Sections III B and IV B should carry over to realistic observations, with a possible caveat concerning nongaussian noise. At this stage, only Viterbi tracking has been tested systematically on spin wandering signals, successfully tracking sources with $h_0 \gtrsim 8 \times 10^{-26}$ and $\varepsilon_{f_\star}(t) < 5 \times 10^{-7}$ Hz for noise levels representative of Advanced LIGO [$S_n(2f_\star)^{1/2} = 4 \times 10^{-24}$ Hz $^{-1/2}$]. CrossCorr and TwoSpect [74] are also expected to track spin wandering sources well, but systematic testing is still underway.

VII. ACKNOWLEDGEMENTS

We would like to thank Paul Lasky, Chris Messenger, Keith Riles, Karl Wette, Letizia Sammut, John Whelan and the LIGO Scientific Collaboration Continuous Wave Working Group for detailed comments and informative discussions. The synthetic data for Stage I of the Sco X-1 MDC were prepared primarily by Chris Messenger with the assistance of members of the MDC team [54]. We thank Chris Messenger and Paul Lasky for their assistance in handling the MDC data. L. Sun is supported by an Australian Postgraduate Award. The research was supported by Australian Research Council (ARC) Discovery Project DP110103347.

[1] K. Riles, Progress in Particle and Nuclear Physics **68**, 1 (2013).

[2] G. Ushomirsky, C. Cutler, and L. Bildsten, MNRAS **319**, 902 (2000).

Index	f_* (Hz)	$h_0(10^{-25})$	$h_0/\sqrt{f_*}$ (10^{-25} Hz $^{-1/2}$)	a_0 (sec)	Estimated f_* (Hz)	Error in f_* (Hz)	Estimated a_0 (sec)	Error in a_0 (sec)	Detection Stage
1	54.498391348174	4.160	0.563524	1.379519	54.4983906268	7.214E-07	1.3795200	1.000E-06	1
2	64.411966012332	4.044	0.503887	1.764606	64.4119658577	1.546E-07	1.7625600	2.046E-03	2
3	73.795580913582	3.565	0.415019	1.534599	73.795580441	4.726E-07	1.5523200	1.772E-02	1
5	93.909518008164	1.250	0.129012	1.520181	93.909517941	6.716E-08	1.5350400	1.486E-02	2
11	154.916883586097	3.089	0.248212	1.392286	154.91689756	1.397E-05	1.3996800	7.394E-03	2
14	183.974917468730	2.044	0.150706	1.509696	183.974917235	2.337E-07	1.0800000	4.297E-01	2
15	191.580343388804	11.764	0.849907	1.518142	191.580298596	4.479E-05	1.5148800	3.262E-03	1
17	213.232194220000	3.473	0.237865	1.310212	213.23225231	5.809E-05	1.3118575	1.646E-03	2
19	233.432565653291	6.031	0.394707	1.231232	233.432550338	1.532E-05	1.2297600	1.472E-03	2
20	244.534697522529	9.710	0.620916	1.284423	244.534696746	7.765E-07	1.2844800	5.700E-05	1
21	254.415047846878	1.815	0.113797	1.072190	254.415047445	4.019E-07	1.0724669	2.769E-04	3
23	271.739907539784	2.968	0.180071	1.442867	271.739936327	2.879E-05	1.4428800	1.300E-05	2
26	300.590450155009	1.419	0.081855	1.258695	300.59044964	5.150E-07	1.0800000	1.787E-01	2
29	330.590357652653	4.275	0.235096	1.330696	330.590357047	6.057E-07	1.3305600	1.360E-04	2
32	362.990820993568	10.038	0.526853	1.611093	362.990070589	7.504E-04	1.5926400	1.845E-02	1
35	394.685589797695	16.402	0.825579	1.313759	394.685618617	2.882E-05	1.3132800	4.790E-04	1
36	402.721233789014	3.864	0.192559	1.254840	402.721233202	5.870E-07	1.2556800	8.400E-04	2
41	454.865249156175	1.562	0.073240	1.465778	454.844343743	2.091E-02	1.4661565	3.785E-04	
44	483.519617972096	2.237	0.101736	1.552208	483.519690961	7.299E-05	1.4601600	9.205E-02	2
47	514.568399601819	4.883	0.215277	1.140205	514.568399349	2.528E-07	1.1404800	2.750E-04	2
48	520.177348201609	1.813	0.079492	1.336686	520.327843196	1.505E-01	1.3370312	3.452E-04	
50	542.952477491471	1.093	0.046897	1.119149	542.952491933	1.444E-05	1.1194380	2.890E-04	3
51	552.120598886904	9.146	0.389254	1.327828	552.120598435	4.519E-07	1.1431103	1.847E-01	2
52	560.755048768919	2.786	0.117639	1.792140	560.755063137	1.437E-05	1.7926028	4.628E-04	3
54	593.663030872532	1.518	0.062283	1.612757	593.663030722	1.505E-07	1.6131735	4.165E-04	3
57	622.605388362863	1.577	0.063198	1.513291	622.56610884	3.928E-02	1.5136818	3.908E-04	
58	641.491604906276	3.416	0.134884	1.584428	641.491619251	1.434E-05	1.5848371	4.091E-04	3
59	650.344230698489	8.835	0.346437	1.677112	650.344215312	1.539E-05	1.6761600	9.520E-04	1
60	664.611446618250	2.961	0.114843	1.582620	664.611402246	4.437E-05	1.5840000	1.380E-03	2
61	674.711567789201	6.064	0.233463	1.499368	674.711611744	4.395E-05	1.5004800	1.112E-03	2
62	683.436210983289	10.737	0.410728	1.269511	683.436269138	5.815E-05	1.2700800	5.690E-04	1
63	690.534687981171	1.119	0.042584	1.518244	690.154763901	3.799E-01	1.5186360	3.920E-04	
64	700.866836291234	1.600	0.060419	1.399926	701.313419622	4.466E-01	1.4002875	3.615E-04	
65	713.378001688688	8.474	0.317256	1.145769	713.373737884	4.264E-03	1.0800000	6.577E-02	1
66	731.006818153273	9.312	0.344417	1.321791	731.006832222	1.407E-05	1.3219200	1.290E-04	1
67	744.255707971300	4.580	0.167871	1.677736	744.254311362	1.397E-03	1.0803545	5.974E-01	2
68	754.435956775916	3.696	0.134556	1.413891	754.435956631	1.449E-07	1.0800917	3.338E-01	2
69	761.538797037770	2.889	0.104699	1.626130	761.720204916	1.814E-01	1.6265499	4.199E-04	
71	804.231717847467	2.923	0.103056	1.652034	804.231732078	1.423E-05	1.6524606	4.266E-04	3
72	812.280741438401	1.248	0.043792	1.196485	812.838541152	5.578E-01	1.1967940	3.090E-04	
73	824.988633484129	2.444	0.085089	1.417154	825.182391835	1.938E-01	1.4175199	3.659E-04	
75	862.398935287248	7.678	0.261467	1.567026	862.398964120	2.883E-05	1.1982204	3.688E-01	1
76	882.747979842807	3.260	0.109728	1.462487	882.74799427	1.443E-05	1.0796966	3.828E-01	2
79	931.006000308958	4.681	0.153408	1.491706	931.006764506	7.642E-04	1.4832000	8.506E-03	2
83	1081.398956458276	5.925	0.180165	1.198541	1081.39898556	2.910E-05	1.1980800	4.610E-04	2
84	1100.906018344283	11.609	0.349877	1.589716	1100.9060324	1.406E-05	1.2086045	3.811E-01	1
85	1111.576831848269	4.553	0.136553	1.344790	1111.57684611	1.426E-05	1.0800000	2.648E-01	2
90	1193.191890630547	0.684	0.019802	1.575127	1193.82006025	6.282E-01	1.5755337	4.067E-04	
95	1324.567365220908	4.293	0.117966	1.591685	1324.56727666	8.856E-05	1.5926400	9.550E-04	2
98	1372.042154535880	5.404	0.145894	1.315096	1372.04216902	1.448E-05	1.0799668	2.351E-01	2

TABLE VI. Results of Viterbi tracking of the 50 closed signals in the Sco X-1 Mock Data Challenge, sorted by index from Ref. [54]. The last column shows in which stage the injected signal is detected.

- [3] A. Melatos and D. J. B. Payne, *Astrophys. J.* **623**, 1044 (2005).
- [4] B. Owen, L. Lindblom, C. Cutler, B. Schutz, A. Vecchio, and N. Andersson, *Physical Review D* **58**, 084020 (1998).
- [5] R. Bondarescu, S. A. Teukolsky, and I. Wasserman, *Physical Review D* **79**, 104003 (2009).
- [6] A. Melatos, J. A. Douglass, and T. P. Simula, *Astrophys. J.* **807**, 132 (2015).
- [7] P. Jaranowski, A. Królak, and B. F. Schutz, *Physical Review D* **58**, 063001 (1998).
- [8] D. I. Jones, *Monthly Notices of the Royal Astronomical Society* **453**, 53 (2015).

- [9] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, *Physical Review D* **77**, 082001 (2008).
- [10] C. T. Y. Chung, A. Melatos, B. Krishnan, and J. T. Whelan, *MNRAS* **414**, 2650 (2011).
- [11] G. Mendell and M. Landry, LIGO Report T050003 (January 2005).
- [12] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, *Physical Review D* **70**, 082001 (2004).
- [13] J. Aasi *et al.*, *Physical Review D* **87**, 042001 (2013).
- [14] J. Aasi *et al.*, *Physical Review D* **93**, 042007 (2016).
- [15] V. Dergachev, LIGO Report T050186 (September 2005).
- [16] V. Dergachev and K. Riles, LIGO Report T050187 (September 2005).
- [17] V. Dergachev, LIGO Report T1000272 (February 2011).
- [18] B. Abbott *et al.*, *Physical Review Letters* **102**, 111102 (2009).
- [19] J. Abadie *et al.*, *Physical Review D* **85**, 022001 (2012).
- [20] J. Aasi *et al.*, *Physical Review D* **93**, 042006 (2016).
- [21] E. Goetz and K. Riles, *Classical and Quantum Gravity* **28**, 215006 (2011).
- [22] J. Aasi *et al.*, *Physical Review D* **90**, 062010 (2014).
- [23] G. Hobbs, A. G. Lyne, and M. Kramer, *Monthly Notices of the Royal Astronomical Society* **402**, 1027 (2010).
- [24] R. M. Shannon and J. M. Cordes, *Astrophys. J.* **725**, 1607 (2010).
- [25] G. Ashton, D. Jones, and R. Prix, *Physical Review D* **91**, 062009 (2015).
- [26] J. M. Cordes and D. J. Helfand, *Astrophys. J.* **239**, 640 (1980).
- [27] S. Price, B. Link, S. N. Shore, and D. J. Nice, *Monthly Notices of the Royal Astronomical Society* **426**, 2507 (2012).
- [28] A. Lyne, G. Hobbs, M. Kramer, I. Stairs, and B. Stappers, *Science* **329**, 408 (2010).
- [29] M. A. Alpar, R. Nandkumar, and D. Pines, *Astrophys. J.* **311**, 197 (1986).
- [30] P. Jones, *Monthly Notices of the Royal Astronomical Society* **246** (1990).
- [31] A. Melatos and B. Link, *Monthly Notices of the Royal Astronomical Society* **437**, 21 (2014).
- [32] J. M. Cordes and G. S. Downs, *The Astrophysical Journal Supplement Series* **59**, 343 (1985).
- [33] G. H. Janssen and B. W. Stappers, *Astronomy and Astrophysics* **457**, 611 (2006).
- [34] K. S. Cheng, *The Astrophysical Journal* **321**, 799 (1987).
- [35] K. S. Cheng, *The Astrophysical Journal* **321**, 805 (1987).
- [36] J. O. Urama, B. Link, and J. M. Weisberg, *Monthly Notices of the Royal Astronomical Society: Letters* **370**, L76 (2006).
- [37] M. deKool and U. Anzer, *Monthly Notices of the Royal Astronomical Society (ISSN 0035-8711)* **262**, 726 (1993).
- [38] A. Baykal and H. Oegelman, *Astronomy and Astrophysics (ISSN 0004-6361)* **267**, 119 (1993).
- [39] L. Bildsten, D. Chakrabarty, J. Chiu, M. H. Finger, D. T. Koh, R. W. Nelson, T. A. Prince, B. C. Rubin, D. M. Scott, M. Stollberg, B. A. Vaughan, C. A. Wilson, and R. B. Wilson, *The Astrophysical Journal Supplement Series* **113**, 367 (1997).
- [40] R. E. Taam and B. A. Fryxell, *The Astrophysical Journal* **327**, L73 (1988).
- [41] A. Baykal, A. Alpar, and U. Kiziloglu, *Astronomy and Astrophysics (ISSN 0004-6361)* **252**, 664 (1991).
- [42] M. M. Romanova, G. V. Ustyugova, A. V. Koldoba, and R. V. E. Lovelace, *The Astrophysical Journal* **616**, L151 (2004).
- [43] A. Baykal, *Astronomy and Astrophysics* (1997).
- [44] B. Abbott *et al.* (LIGO Scientific Collaboration), *ApJL* **683**, L45 (2008).
- [45] B. G. Quinn and E. J. Hannan, *The Estimation and Tracking of Frequency* (Cambridge University Press, 2001) p. 266.
- [46] S. Paris and C. Jauffret, *IEEE Transactions on Aerospace and Electronic Systems* **39**, 439 (2003).
- [47] L. White and R. Elliott, *IEEE Transactions on Signal Processing* **50**, 1205 (2002).
- [48] J. J. Williams and A. K. Katsaggelos, *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council* **13**, 900 (2002).
- [49] R. Streit and R. Barrett, *IEEE Transactions on Acoustics, Speech, and Signal Processing* **38**, 586 (1990).
- [50] R. Barrett and D. Holdsworth, *IEEE Transactions on Signal Processing* **41**, 2965 (1993).
- [51] X. Xie and R. Evans, *IEEE Transactions on Signal Processing* **39**, 2659 (1991).
- [52] X. Xie and R. Evans, *IEEE Transactions on Signal Processing* **41**, 1391 (1993).
- [53] A. Viterbi, *IEEE Transactions on Information Theory* **13**, 260 (1967).
- [54] C. Messenger, H. Bulten, S. Crowder, V. Dergachev, D. Galloway, E. Goetz, R. Jonker, P. Lasky, G. Meadors, A. Melatos, S. Premachandra, K. Riles, L. Sammut, E. Thrane, J. Whelan, and Y. Zhang, *Physical Review D* **92**, 023006 (2015).
- [55] L. Sammut, C. Messenger, A. Melatos, and B. Owen, *Physical Review D* **89**, 043001 (2014).
- [56] J. Aasi *et al.*, *Physical Review D* **91**, 062008 (2015).
- [57] R. Prix, LIGO Report T0900149 (June 2011).
- [58] A. Melatos, C. Peralta, and J. S. B. Wyithe, *The Astrophysical Journal* **672**, 1103 (2008).
- [59] C. Espinoza, A. Lyne, B. Stappers, M. Kramer, M. Burgay, N. Dmico, P. Esposito, A. Pellizzoni, and A. Posenti, in *RADIO PULSARS: AN ASTROPHYSICAL KEY TO UNLOCK THE SECRETS OF THE UNIVERSE. AIP Conference Proceedings*, Vol. 1357 (2011) pp. 117–120.
- [60] R. Bellman, *Princeton University Press Princeton New Jersey*, Vol. 70 (1957) p. 342.
- [61] R. Prix, *Physical Review D* **75**, 023004 (2007).
- [62] J. Aasi *et al.*, *Classical and Quantum Gravity* **32**, 074001 (2015).
- [63] D. Rife and R. Boorstyn, *IEEE Transactions on Information Theory* **20**, 591 (1974).
- [64] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables* (Courier Corporation, 1964).
- [65] D. K. Galloway, S. Premachandra, D. Steeghs, T. Marsh, J. Casares, and R. Cornelisse, *The Astrophysical Journal* **781**, 14 (2014).
- [66] J. T. Whelan, S. Sundaresan, Y. Zhang, and P. Peiris, *Physical Review D* **91**, 102005 (2015).
- [67] S. W. Ballmer, *Classical and Quantum Gravity* **23**, S179 (2006).
- [68] B. Abbott *et al.*, *Physical Review D* **76**, 082003 (2007).
- [69] J. Abadie *et al.*, *Physical review letters* **107**, 271102 (2011).

- [70] C. Messenger and G. Woan, Classical and Quantum Gravity **24**, S469 (2007).
- [71] S. van der Putten, H. J. Bulten, J. F. J. van den Brand, and M. Holtrop, Journal of Physics: Conference Series **228**, 012005 (2010).
- [72] J. Aasi *et al.*, Physical Review D **91**, 062008 (2015).
- [73] A. L. Watts, B. Krishnan, L. Bildsten, and B. F. Schutz, Monthly Notices of the Royal Astronomical Society **389**, 839 (2008).
- [74] G. D. Meadors, E. Goetz, and K. Riles, Classical and Quantum Gravity **33**, 105017 (2016).